

Carnegie Mellon
Software Engineering Institute

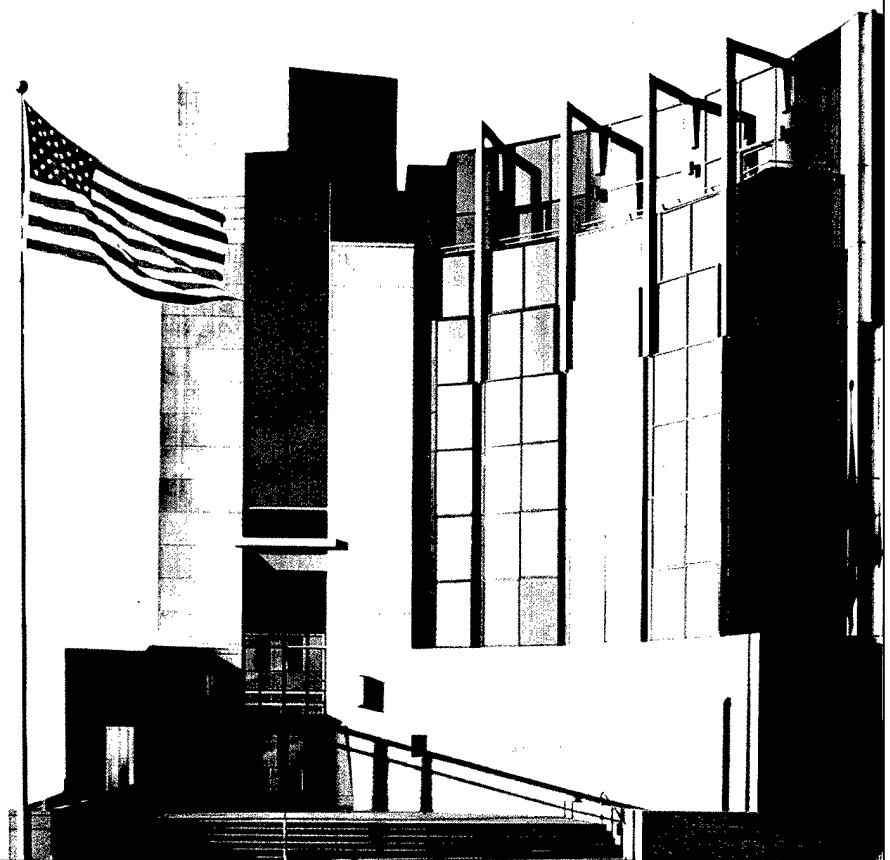
A Simulation Model for Managing Survivability of Networked Information Systems

Soumyo D. Moitra
Suresh L. Konda

December 2000

TECHNICAL REPORT
CMU/SEI-2000-TR-020
ESC-TR-2000-020

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited



Carnegie Mellon University does not discriminate and Carnegie Mellon University is required not to discriminate in admission, employment, or administration of its programs or activities on the basis of race, color, national origin, sex or handicap in violation of Title VI of the Civil Rights Act of 1964, Title IX of the Educational Amendments of 1972 and Section 504 of the Rehabilitation Act of 1973 or other federal, state, or local laws or executive orders.

In addition, Carnegie Mellon University does not discriminate in admission, employment or administration of its programs on the basis of religion, creed, ancestry, belief, age, veteran status, sexual orientation or in violation of federal, state, or local laws or executive orders. However, in the judgment of the Carnegie Mellon Human Relations Commission, the Department of Defense policy of "Don't ask, don't tell, don't pursue" excludes openly gay, lesbian and bisexual students from receiving ROTC scholarships or serving in the military. Nevertheless, all ROTC classes at Carnegie Mellon University are available to all students.

Inquiries concerning application of these statements should be directed to the Provost, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, telephone (412) 268-6684 or the Vice President for Enrollment, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, telephone (412) 268-2056.

Obtain general information about Carnegie Mellon University by calling (412) 268-2000.



Carnegie Mellon
Software Engineering Institute
Pittsburgh, PA 15213-3890

A Simulation Model for Managing Survivability of Networked Information Systems

CMU/SEI-2000-TR-020
ESC-TR-2000-020

Soumyo D. Moitra
Suresh L. Konda

December 2000

Networked Survivable Systems

Unlimited distribution subject to the copyright.

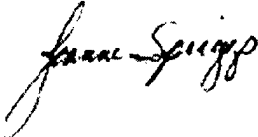
20010420 020

This report was prepared for the

SEI Joint Program Office
HQ ESC/DIB
5 Eglin Street
Hanscom AFB, MA 01731-2116

The ideas and findings in this report should not be construed as an official DoD position. It is published in the interest of scientific and technical information exchange.

FOR THE COMMANDER



Joanne E. Spriggs
Contracting Office Representative

This work is sponsored by the U.S. Department of Defense. The Software Engineering Institute is a federally funded research and development center sponsored by the U.S. Department of Defense.

Copyright 2000 by Carnegie Mellon University.

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

Internal use. Permission to reproduce this document and to prepare derivative works from this document for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use. Requests for permission to reproduce this document or prepare derivative works of this document for external and commercial use should be addressed to the SEI Licensing Agent.

This work was created in the performance of Federal Government Contract Number F19628-00-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 52.227-7013.

For information about purchasing paper copies of SEI reports, please visit the publications portion of our Web site (<http://www.sei.cmu.edu/publications/pubweb.html>).

Table of Contents

Acknowledgements	vii
Abstract	ix
1 Introduction	1
1.1 The Problem Area	1
1.2 Importance/Motivation	1
1.3 Literature Review and Taxonomy	2
1.4 Objectives	2
1.5 Approach	3
1.6 Outline	4
2 Model Development	5
2.1 Notation	5
2.2 Model for the Incidents Process	5
2.3 Model for the System	7
2.4 Survivability Modeling	10
3 Simulation of Episodes	15
4 Simulation Results and Analysis	17
5 Summary and Conclusions	25
5.1 Summary	25
5.2 Conclusions	25
5.3 Policy Implications	26
6 Future Work	29
Bibliography	31

Blank Page.

List of Figures

Figure 1: Components of the Simulation Model	3
Figure 2: The Marked Stochastic Point Process	6
Figure 3: Expected Survivability Against Cost	23

Blank Page.

List of Tables

Table 1:	Expected Survivability and $P(j)$	17
Table 2:	Expected Survivability and π_1	18
Table 3:	Expected Survivability and χ_1	18
Table 4:	Expected Survivability and π_3	19
Table 5:	Expected Survivability and χ_3	19
Table 6:	Average Damage and Arrival Rate a	20
Table 7:	Expected Survivability and Correlation of (a, j)	21
Table 8:	Relative Changes in Survivability with Respect to Parameter Values	21

Acknowledgements

This work was done while Soumyo Moitra was a Visiting Scientist at CERT[®]. We wish to acknowledge the support and encouragement of Tom Longstaff. We also wish to thank the members of the support staff at CERT, particularly Stephanie Rieger and Annette Welsch for their help.

Moitra Soumyo

Suresh L. Konda

December 2000

[®] CERT and CERT Coordination Center are registered in the U.S. Patent and Trademark Office.

Abstract

In this paper we develop a model to evaluate the tradeoffs between the cost of defense mechanisms for networked systems and the resulting expected survivability after a network attack. The model consists of three submodels. The first submodel simulates the occurrence of attacks or incidents. The second submodel simulates the impact of an attack on the system. This depends on the type of attack and the defense mechanism installed in the system. The third submodel assesses the survivability of the system, which depends on the degree of its degradation after the attack.

By varying the level of defense in the simulation, we examine how this expected survivability changes with the defense level. Since costs are assumed to increase with the strength of the defense system, we can derive a cost/survivability curve that managers can use to decide on the appropriate level of security for their organizations. We have also explored the sensitivity of expected survivability to various parameters of the model, such as the mix of attack types and the rate of occurrence of incidents.

1 Introduction

1.1 The Problem Area

Today's information systems are increasingly linked together in an unbounded network [Ellison 97, Fisher 99]. At the same time two major trends are occurring. One is that individuals and organizations are becoming critically dependent on computer networks, and two, the vulnerability of a network system is increasing since far more potential attackers are having access to the network and hence to other people's systems.

It appears that malicious attacks of various kinds will inevitably occur in unbounded networks such as the Internet, and it also seems inevitable that some of these attacks will cause damage to systems and loss to their owners. The damage done by and costs of such attacks have been estimated at varying levels, but it is clear that even by conservative estimates, they are considerable [CSI 98]. Therefore it has become necessary for systems managers and researchers to find ways to improve the security of Information Systems (ISs).

However, there is probably no **absolute** security, and the real issue is the level to which we deploy defense mechanisms against these random attacks in general. That is, we need to enhance security, but we need to decide by *how much* to enhance it—given costs. In other words, we have to determine how to enhance network security for ISs *efficiently*. We would like to achieve the “optimal” or the most cost-beneficial level of security based on our needs, financial abilities, and potential threats.

1.2 Importance/Motivation

In view of the above discussion on the dangers and costs of network attacks in unbounded network environments and the Internet, **cost/benefit analysis** to enhance security efficiently becomes important. The costs will be those of deploying various defense mechanisms to protect a system/site against attacks. The benefits will be those of increased survivability of the system/site. Survivability means the ability of systems to recover from attacks, and in particular the degree to which they recover [Ellison 97]. Thus we need to explore methods to improve the survivability of network systems in cost-beneficial ways.

One approach that would help achieve this is to model the occurrence of attacks, model system-response and the impact of attacks, and simulate alternative scenarios to examine how different parameters affect system survivability. In this paper, we develop such a set of models and simulate them to analyze network survivability under various conditions.

1.3 Literature Review and Taxonomy

There has been considerable work done on survivability in Telecommunications, but that is based on topological considerations, where the impact of link or node failures are studied [Moitra 97]. There is some literature in this area as far as network information systems are concerned [Howard 95, Fisher 99, Ellison 97, Linger 98]. One of the key issues is the taxonomy that is to be used in discussing network security incidents and survivability. Many alternative terms have been used, and here we follow the attempt to develop a common language for computer security incidents [Howard 98]. This language defines a number of terms, in particular, an attack and an incident. An attack is defined as

“a series of intentional steps taken by an attacker to achieve an unauthorized result.”

An incident is defined as

“a group of related attacks that can be distinguished from other attacks because of the distinctiveness of the attackers, attacks, objectives, sites, and timing.”

The literature also classifies terms such as tools, actions, targets, and results. All these terms are important since they clarify concepts that are required to study and understand security and survivability of systems. For this paper, it would be useful to add some additional terms. We define a *foray* as a combination of [tool + action + target] and an *episode* as the combination [incident + response], that is the whole process of a set of attacks and the system's response to the incident as a whole. In this paper, we consider a *system* to be the collection of all the relevant computers and network elements at a site. Finally we will refer to the system's *configuration* as the combination of its design and its defense mechanism.

1.4 Objectives

A primary objective of this paper is to develop a simulation model that can be used by systems managers and chief information officers (CIOs) to understand survivability issues and evaluate the tradeoffs involved in decisions about network systems design and defense mechanisms. Along with this, we wish to

- suggest data collection and analysis strategies for understanding and forecasting patterns of attacks on sites
- derive measures of survivability of network and information systems
- set up a methodological framework for cost/benefit analysis of security in systems
- provide heuristics for moving towards “optimal” or improved security strategies
- develop a simulation model that could be a basis for a Decision Support System (DSS) to manage systems security and survivability

1.5 Approach

Our approach is illustrated in Figure 1. The goal is to assess the survivability of a system when it is subjected to a series of random incidents over time, where incidents are as defined above. For this reason, we first need to model the process of occurrence of incidents from the point of view of a system or site that experiences this process over time. This is equivalent to a stochastic point process where incidents occur at random points in time; therefore we need to simulate a stochastic point process. The survivability also depends on how the system responds to an incident. This will depend on the system configuration, that is, its design and defense mechanisms as defined above. Therefore, we need to model this response as a function of the incident type and configuration. The model will involve a transition matrix that will give the probabilities of the system ending up in any of its possible states after experiencing an incident. These probabilities will depend on the incident type and system configuration. Next, the degree to which it has survived will have to be measured. This will be a function of the state in which it ends up and the amount of compromise that has occurred. For this purpose, we develop some new survivability measures that take into account the different dimensions of survivability, that is, the different functionalities and services that can be compromised.

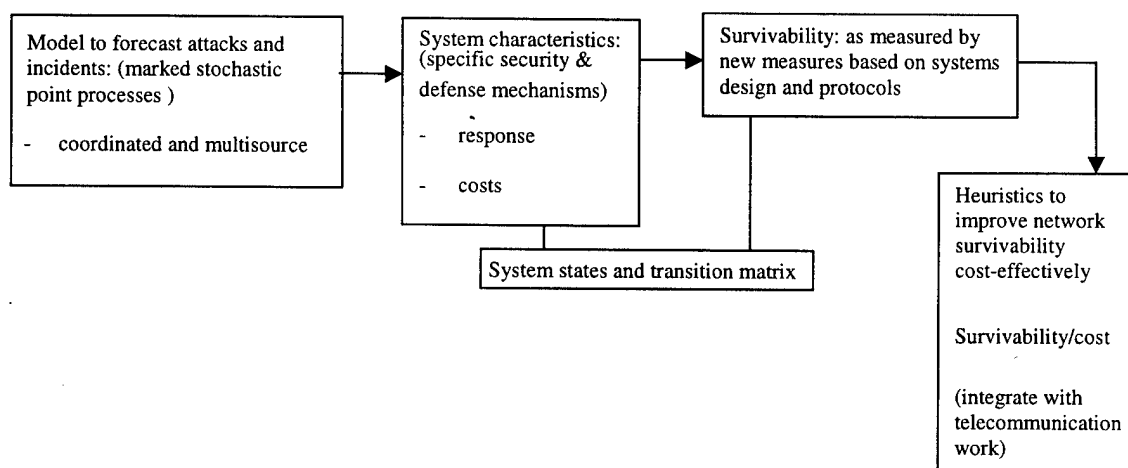


Figure 1: Components of the Simulation Model

This view is depicted in Figure 1. With this simulation model we can analyze the costs and benefits of alternative defense mechanisms under various scenarios. Based on such analyses, systems managers can make decisions regarding the systems configuration that best suits their needs.

The advantage of this systems simulation approach is that a large variety of scenarios may be explored. Alternative incident processes, different systems configurations, various state transition probabilities, and additional survivability measures may all be investigated with such a model. Thus, given the high degree of uncertainty regarding future attacks and their impacts, this method provides a practical approach to assess and manage survivability.

1.6 Outline

The remainder of the paper is organized as follows. Section 2 develops the model that we shall use. Section 3 describes the simulation procedure and assumptions. The results of the analysis are given in Section 4. The implications of the results are discussed in Section 5, and Section 6 outlines some of the future work that should be done.

2 Model Development

As illustrated in Figure 1, we shall develop models of the incidents-process, the response of the system, and its survivability. We shall use the following notation. Additional notation will be introduced as needed.

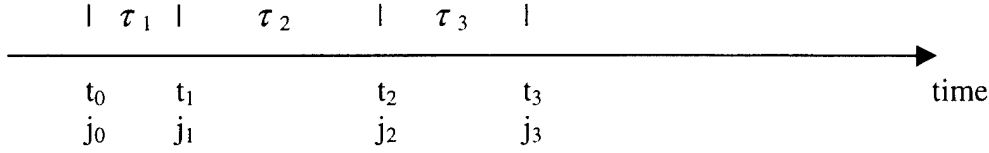
2.1 Notation

- i] i, j = index for incident type, i, j in $\{J\}$. We consider actual, unauthorized incidents only. i denotes the prior incident and j the subsequent (or current) one.
- ii] $P(j)$ = probability that an incident is of type j .
- iii] $\tau(i, j)$ = inter-incident times between incidents i and j .
- iv] a = arrival rate of incidents = $1/\tau$.
- v] r, s = index for system state, r, s in $\{S\}$.
- vi] d = index for system design, d in design space $\{D\}$.
- vii] m = index for defense mechanism, m in the set $\{M\}$.
- viii] configuration = design \times mechanism in configuration space $\{D \times M\}$.
- ix] T = transition probability matrix with elements $\{p(r, s)\}$, where $\{p(r, s)\}$ possibly being functions of i, j, d, m .
- x] l = (victim) sites, l in $\{L\}$.
- xi] $h(l)$ = index for incidents at individual site l : $h(l) = 1, 2, 3, \dots$
- xii] $H(l)$ = total number of incidents at site l .
- xiii] $t(h(l), l)$ = time of h -th incident at l
$$= \sum_{k=1}^{h(l)} \tau(k), \text{ where } \tau(k) = t(k) - t(k-1).$$
- xiv] n = number of simultaneous attacking sites in an incident.
- xv] $g(n | v)$ = probability density function for n with parameter v .

2.2 Model for the Incidents Process

In order to forecast incidents, we model the process as a marked, stochastic point process, where the incidents are the events that occur at random points in time, and the event type is the mark associated with an incident [Snyder 91]. The mark is used to identify random quantities associated with the point it accompanies. As shown in Figure 2, each occurrence time t_k of the k -th incident in a temporal point-process has a mark j_k associated with it, where j_k will have values in a specified space. The mark, or event type in our case, has to take into account the severity of the incident and the possibility of single, or multiple and simultaneous attacks. This is because we are modeling a process that is taking place in an unbounded environment [Ellison 97]. Therefore the mark space will be 2-dimensional, characterized by type (severity) and number-of-attackers. That is, it will be in the $\{J \times N\}$ space. Although this 2-D marked

point process model was developed, no data on the distribution of the number of attackers per incident were available, so only a 1-D mark space with severity was used in the simulations.



$\tau \sim$ inter-incident time;
 $t \sim$ times at which incidents occur;
 $j \sim$ marks associated with each incident (incident type).

Figure 2: The Marked Stochastic Point Process

A stochastic point process can generally be represented as $\{x(t): t \in T\}$, that is, as a family of random variables indexed by a parameter t that takes values in a parameter set T called the index set of the process. In our case, t represents time, and since T is a subset of \mathbb{R} , it is a continuous-parameter process. The stochastic point process $\{x(t): t \in T\}$ is completely characterized statistically by the joint distribution function

$$P_{x(t_1), x(t_2), \dots, x(t_k)}(X_1, X_2, \dots, X_k) = \Pr(x(t_1) \leq X_1, x(t_2) \leq X_2, \dots, x(t_k) \leq X_k)$$

for the random variables $x(t_1), x(t_2), \dots, x(t_k)$ known for any finite collections (t_1, t_2, \dots, t_k) and $\{X_1, X_2, \dots, X_k\}$ where $t_i \in T$ and $X_i \in \mathbb{R}$ for $i = 1, 2, \dots, k$. With every point process, there is an associated counting process denoted by $\{N(t): t \geq t_0\}$ which indicates the total number of points in the interval $[t_0, t)$ regardless of their marks [Snyder 91].

The characteristic functional for a sequence of independent random variables is given by

$$\Phi_{x(i \nu)} = E[\exp\{i \int_{t_0}^T \nu'(t) dx(t)\}]$$

where $\{\nu(t): t_0 \leq t \leq T\}$ is an arbitrary vector-valued function, the prime denotes the transpose operation, and $i = \sqrt{-1}$ here. For the purposes of this analysis, we limit our attention to the probability density function of the “inter-incident times” (τ ’s) which we denote by $f(t)$. That is,

$$f(t) = \Pr\{t \leq \tau \leq t + dt\}.$$

When the process is Poisson, the density function is given by

$$f(t) = a * e^{-at}$$

where a is the rate of occurrence of incidents, and the distribution function is given by

$$F(t) = 1 - e^{-at}.$$

For this paper, we use hypothetical data to run the simulations, but they are based on actual records, and the model is also based on observations of actual events. For future applications, the parameters of the model developed here should be estimated from actual data on incidents. There are a number of issues in estimating the model parameters for the incidents-process:

1. The functional forms for inter-event times, $f(t)$ must be determined. Frequently this is assumed to be Poisson [Law 82], but this has to be verified by examining the distributions observed from the data. It may be that some other distribution such as the Weibull, or a mixture of exponential distributions will be more appropriate.
2. Once the form for $f(t)$ has been determined, its parameters will have to be estimated [Basawa 80].
3. Next the $P(j)$'s, or the probabilities of each incident type j will have to be estimated [Basawa 80].
4. It is also important to test whether $f(t)$ depends on i , or j , or both.
5. Similarly, the issue of stationarity of the process will have to be investigated.
6. Finally, it is necessary to check whether the incidents process depends on victim site (for example its domain type), and other dependencies in general.

We should note here that the incidents recorded in actual data are likely to be twice filtered: that is, they are conditional upon detection, and then, upon reporting. Also, any recorded data will be doubly-censored data, that is, both right and left censored. This means that the process had already started before data collection began, and the process had not finished when data collection was stopped (at least for most realistic data sets). Censoring may introduce biases in parameter estimates, and it is important to take note of this.

2.3 Model for the System

Next we need to characterize the systems designs under consideration and the potential defense mechanisms that may be employed within the systems. That is, we need to define the design/architecture space $\{D\}$ of the system, and the defense mechanism state space $\{M\}$. The combination of a system design and defense mechanism will be called the configuration (or posture) space, $\{D \times M\}$. The design could include distributed sub-systems with different defenses for the sub-systems. Each possible alternative would be a configuration. Initially we assume one design only. When information exists on different designs, any number of designs may be analyzed. We also assume five hypothetical levels of defense mechanisms, and cost increasing with effectiveness. In general, many complex designs and defense mechanisms can exist, and our model can accommodate such complexity whenever the data are available.

The response prediction model will predict the transition of the system to a new state after an attack/incident has occurred, and will be a function of the incident type and the configuration, or $p(r,s) = p(r,s | j,d,m)$. Thus, given an incident-type j and initial system state r , the subsequent state s may be any one of the set $\{S\}$ of possible states that the system can be in, such as normal, under attack, compromised, recovered, or non-functional. The actual states may of course be different for different configurations. The transition matrix T will probabilistically map r to s given j, d, m . That is, each element of T is the probability of the system of design d and defense mechanism m going to another (possibly compromised) state when subjected to an incident of type j . In general, the incident type j will be a vector of severity level and number of attackers. But since data on the number of attackers were not available, j is taken to be severity only in the simulations conducted here.

We assume the following structure for T . Without any loss of generality, we assume that the states are ordered by degree of compromise, that is, from $s = 1 = \text{normal}$ (totally functioning) to $s = S = \text{(totally) non-functional}$. Given an incident, the system can never go to a “better” state: therefore the lower triangle below the diagonals will have structural zeros as shown below.

$$\begin{pmatrix} p_{11} & p_{12} & p_{13} & p_{14} & p_{15} \\ 0 & p_{22} & p_{23} & p_{24} & p_{25} \\ 0 & 0 & p_{33} & p_{34} & p_{35} \\ 0 & 0 & 0 & p_{44} & p_{45} \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

We also impose the following constraints on the elements of T , $\{p(r,s)\}$, in terms of their dependence on s, j , and m .

$p(r,s) \downarrow s, \forall s > r$, holding j, m constant, that is, same severity level and same defense;

this implies graceful degradation: the probability of going to a much worse state is lower than going to a slightly worse state.

$p(r,s) \uparrow r, \forall s > r$, holding j, m constant, that is, same severity level and same defense;

vulnerability increases with level of degradation.

Assuming that the j 's are ordered from most severe to least severe,

$p(1,1) \uparrow j$, holding m constant, that is, same defense level;

probability of staying normal is higher if the incident is less severe.

$p(1,s) \downarrow j, \forall s > 1$, holding m constant, that is, same defense level;

probability of degradation is lower if the incident is less severe.

$p(r,s) \downarrow m, \forall s > r$, holding j constant, that is, same severity level;

probability of degradation is lower if the defense is stronger.

$p(r,r) \uparrow m, \forall r$, holding j constant, that is, same severity level;

probability of staying in the same state and not degrading is higher if the defense is stronger.

$p(r,s) \uparrow n, \forall s > r$, holding all else constant;

probability of degradation increases with the number of attackers.

And $\sum_s p(r,s) = 1, \forall r$. This implies that the system must end up in some state or other.

If we know the transition probabilities in each case, we can input that data directly into the model. Otherwise, we can develop a model to generate the elements $\{p(r,s)\}$ of the transition probability matrix T , or compute them by considering the path through intermediate states during the attack-response episodes that the system may experience. Estimating these transition matrices is critical but extremely complex, since $S^2 \times J \times D \times M$ probabilities must be estimated. Thus some simplifying rules may be employed to generate them. Since data on these probabilities were not available, we considered only one value of D ; that is, we assumed only one design in the simulation runs below. The model, however, has been designed to handle any number of designs as long as data on them are available.

Currently, no reliable data are available on the times to transition to different states, or the time to fully recover. The CERT[®] data indicate that the mean time between incidents *at a site* is greater than one month. Since it may be reasonably expected that recovery times will be shorter than that on the average [Cohen 98], in these simulations we have assumed that the system would always fully recover before the next incident occurred. So the initial state r was always set equal to 1. However, the model includes the possibility of the system still being in a compromised state when the next incident occurs. We can simulate these conditions given data on system transition times.

In the absence of data, we developed a model to generate the $p(1,s)$'s, such that

[®] CERT and CERT Coordination Center are registered in the U.S. Patent and Trademark Office.

$$p(1,s) = p(s, j, \text{cost}(m); \pi_0, \chi_0, \pi_1, \chi_1, \pi_2, \chi_2).$$

There are two cases, $s = 1$ and $s > 1$.

$$p(1,1) = \pi_2 * (1 - e^{-\pi_1^{\text{cost}(m)} - \pi_0}) \quad \text{for } s=1, \text{ and}$$

$$p(1,s) = \chi_2 * (e^{-\chi_1^{\text{cost}(m)} - \chi_0}) \text{ for } s > 1.$$

These are simple but commonly used functional forms that are concave and convex respectively, and so reflect decreasing returns with cost. π_1 and χ_1 are the critical shape coefficients that determine the relationship of the transition probabilities with the cost of the defense mechanisms $\text{cost}(m)$. This in turn determines how the survivability varies with cost.

$\pi_2 = \pi_2(j)$ which is modeled as a linear function $= \pi_3 * j$, and

$\chi_2 = \chi_2(j,s) = \chi_3 * ((6-s) - (.4*j))$, again linear in s and j .

The scale coefficients π_3 and χ_3 as well as the constants were calibrated to give reasonable values of the transition probabilities subject to all the restrictions given above. The location coefficients π_0 and χ_0 were set to 0, and $\pi_1, \chi_1, \pi_3, \chi_3$ were varied during the simulation runs.

In addition, since we can never be certain of the estimates, sensitivity analysis should be performed. For this, we may need to have guidelines on the range of possible variations of these $\{p(r,s | j,d,m)\}$'s. Alternatively, we might be able to estimate confidence intervals (CIs), within which to vary these probabilities.

In the future, with data available, it may be possible to cluster the incident types and/or the system configurations into smaller subsets, since our interest is with respect to their impact on survivability only. This will make it easier for managers to assess the survivability of their systems.

2.4 Survivability Modeling

As discussed in the introduction, survivability is the key issue we wish to investigate with the simulation model. Therefore it is necessary to develop a measurable concept of survivability. There has been considerable work done on survivability in telecommunications [Moitra 97], and although that analysis is essentially at the network topology level, we may be able to arrive at concepts suitable to information systems and networks.

Survivability is the degree to which a system has been able to withstand an attack or attacks, and is still able to function at a certain level in its new state after the attack. This new state s , in general will be a compromised state, and is the state in which the system ends up before

any full-fledged recovery or repairs are done to it to restore it to its normal state. At the conceptual level, we propose that survivability be measured as:

$$\text{SURV} = (\text{performance level at new state } s) / (\text{normal performance level})$$

The main issue is the measurement of performance levels. In telecommunications, it is generally taken as the traffic that is still carried relative to the offered traffic the network could carry under normal conditions. An analogous approach could be taken for computer systems, in that the different functionalities and services could be considered separately, and an assessment could be made as to what extent each functionality has survived in the new system state after an attack. For example, if a given functionality has survived intact, its value would be 1, and if the system were completely nonfunctional with respect to that service, then its value would be 0. Intermediate states would have values in between.

Let $\varphi(s, k)$ be the degree to which the compromised function/service k has survived in state s , and let $w(k)$ be the importance level of function/service. Then one possible measure of survivability might be in the form of a weighted sum:

$$\text{SURV}(s) = \sum_k w(k) * \varphi(s, k)$$

This assumes that a complete set of states $\{S\}$ of the system has been defined, and that a systems analyst or IS manager can assess $\varphi(s, k)$ for each s and k . In view of the data requirements, it may be necessary to aggregate the state space $\{S\}$, and the different functionalities and services $\{K\}$. The states in $\{S\}$ may be {normal, under attack, compromised, recovered, non-functional}, for example, or {normal, minor compromise, significant compromise, very serious compromise, nonfunctional}. Then $\varphi(s, k)$ could be the average level to which function or service k survives in each of those states s . This is a flexible approach, and can be applied in many situations. For example, there might be a particular function that an organization values very highly (such as protecting the confidentiality of a database in a financial services company). Then the weight on this would be very high and also the survivability of this function could be rated low even for a slight compromise. Then any defense mechanism that protected this function would give a high expected survivability, and thus a high benefit, while a defense that did not protect this function would give very low value for expected survivability, and thus very low benefits.

This is a standard multi-criteria approach to assessing survivability. While this approach has been used widely, there can be difficulties and biases associated with such a measure. These can be mostly overcome through careful analysis. The weights $w(k)$ are such that

$$0 \leq w(k) \leq 1, \text{ and } \sum_k [w(k)] = 1;$$

The $\varphi(s,k)$'s may also be normalized measures $0 \leq \varphi(s,k) \leq 1$. Then SURV(s) will be between 0 and 1, where 0 means total failure and 1 means completely normal.

Another measure may be a "relative" survivability measure. To derive this, we consider the maximum level of functionality k in the normal state, $X(k)$. However, the level of functionality *required* might be $x(k)$ where $0 \leq x(k) \leq X(k)$. Let the level of functionality k that has survived in state s be $x'(k,s)$. Then we can define *survivability relative to the requirement* as

$\varphi'(k,s) = x'(k,s)/x(k)$ if $x' < x$, that is, the fraction of the required level that is available, and

$\varphi'(k,s) = 1$ if $x' > x$, since the surviving level is greater than what is required.

Instead of a weighted survivability, we may consider the worst degree of compromise that has occurred across all functions and services. This would be analogous to "worst-case" survivability that is often considered in telecommunications. In that case,

$$\text{SURV}(s) = \min_k \varphi(k,s).$$

In many real situations one cannot be always aware of every possible vulnerability of a system. However, it may be possible to enumerate the set of all possible compromises that *could* occur given the existence of (unknown or known) vulnerabilities. In such cases, we may proceed as follows:

Let the probability that function/service k is compromised to degree x by incident-type j be given by $p_{k,j}(x)$. Then we can simulate the overall compromise across all k , and compute the survivability after each incident, or we can simplify the analysis and consider the *expected* compromise $E[x(k,j)]$ given j ,

$$\text{Where } E[x(k,j)] = \int_0^1 x * p_{k,j}(x) * dx \quad \text{assuming } 0 \leq x \leq 1$$

and compute survivability as

$$\text{SURV} | j = \sum_k w(k) * (1 - E[x(k,j)])$$

It may be desirable that the weights $w(k)$ reflect the utilization of a function or service in addition to its importance. On the other hand, we still need to distinguish the "essential" functions and services regardless of how much they are utilized. In such situations, we might partition the set $\{K\}$ into say, $\{K_0, K_1, K_2\}$, where K_0 is a set of unimportant functions/services, K_1 is a set of functions/services that are used/needed very often, but not critical, and K_2 is a set of essential functions/services.

$$\text{SURV}(s) = \left[\prod_{k'} \varphi(s, k')^{w(k')} \right] * \left[\sum_k w(k) * \varphi(s, k) \right]$$

$$\text{or SURV}(s) = \frac{1}{2} * \left[\prod_{k'} \varphi(s, k')^{w(k')} \right] * \left[1 + \sum_k w(k) * \varphi(s, k) \right]$$

where k is in $\{K_1\}$ and k' is in $\{K_2\}$. The multiplicative term in both cases ensures that if an essential function/service fails, that is $\varphi(s, k') = 0$ for any k' , survivability goes to 0. The second form for $\text{SURV}(s)$ ensures that it does not always go to 0 when all the functions/services in set $\{K_2\}$ fail totally but some or all the functions/services in set $\{K_1\}$ have survived.

Another approach is to consider the relative degree of survivability of different systems and their configurations. It may be possible to conduct pair-wise comparisons of systems and sites and assess relative survivability of one site with respect to the other. If sufficient data can be obtained on some sites, then our proposed simulation model can be run to explore a variety of potential attack scenarios. Yet another approach to assessing survivability is to use "conjoint analysis" which is a method of evaluating the aggregate value of a product or service (in this case survivability) based on "part-worths" of the different feature or aspects that comprise survivability [Lilien 92].

We recognize that there is no "absolute" survivability. Some attack or other may compromise any system, however well defended. What we are interested in is assessing the strength of a current defense mechanism of a system of a given design relative to a stochastic incidents process. The actual survivability could be a function of many other factors such as the policies of the system managers, the "behavior" of the system and the deterrence it can induce among potential attackers, its reaction (detection, resistance, recovery), or the publicity surrounding an incident experienced by the system.

3 Simulation of Episodes

The simulation model that we have developed simulates whole episodes, that is, the process starting from the generation of an incident, through the response of the victim system or site, to its average survivability with a given defense mechanism. This also allows us to do cost/benefit analyses, where cost is the cost of the defense mechanism and the benefit is the increased survivability of the system.

The steps in the simulation are as follows:

Initializations: a = arrival rate (mean inter-event time (IET) $= 1/a$.)
 v = parameter for the distribution of n
 J = total number of incident types
 $P(j)$ = probabilities of each type of incident occurring
 M = total number of defense mechanisms
 $\text{cost}(m)$ = cost of each defense mechanism m
 S = total number of states the system can be in
 $\text{SURV}(s)$ = survivability of the system in each state s .
 READ IN DATA FROM INPUT FILE:

$f(t)$ is set:

Generate an incident: This means randomly generating $(t, j, n) \mid f(t), P(j), g(n)$;

Generate new state of system (r to s) from $T(j, m, d)$

(N.B. another incident may occur before stable state/recovery!
 Can handle that | data on transient times.)

Compute survivability = $\text{SURV}(s)$:

Repeat incidents – till <end condition = TRUE> End condition = #incidents or total time

Average survivability $\mid f, d, m$. This completes one run

Write summary report: $p(j), \text{Cost}(m), \text{EXP. SURV.} \mid a, \pi_1, \chi_1, \pi_3, \chi_3$

Vary m ; repeat for $\{m \text{ in } M\}$ This means a new T :

 Add to summary report .

Plot EXP-SURV|m versus cost(m) given other parameter values.

Perturb parameter values – sensitivity analysis:

Table of cost(m), EXP-SURV; given the new parameter values.

END.

Assumptions

Form for $f(t)$ is exponential; (can easily be changed to another distribution).

$P(j)$'s discrete; given for each incident type j .

Geometric distribution $g(n)$ for n with parameter v .

$J=2$; $M=6$; $D=1$; Cost(m) = hypothetical values - scaled from 1 to 100.

$T=\{p(r,s)\}$; computed from the values of $\pi_0, \pi_1, \pi_2, \pi_3, \chi_0, \chi_1, \chi_2$, and χ_3 .

SURV(s) – assumed values; (to be computed in actual applications).

4 Simulation Results and Analysis

Our interest is to observe how well a system survives when subjected to a series of attacks. This will obviously depend on both the severity levels of the attacks as well as the level of defense that is built into the system. The stronger the defense system, the more likely it is to withstand an attack, that is to stay in its normal state, and less likely to end up in a compromised state. In other words, the transition probabilities of the system are a function of the defense mechanism, and this functional relationship drives the expected survivability of the system in any attack scenario. Therefore simulation was carried out for different probabilities of the attack types, and different relationships between the cost of the defense mechanism and the probabilities of the system ending in the various possible states (from normal to nonfunctional).

A large number of simulations can be carried out with our model to investigate a wide variety of issues related to managing survivability, since we can observe the impact of any model parameter on the system survivability. In this paper we present some of the possible sensitivity analyses to illustrate what can be done. The results are presented in Tables 1 to 7 where the survivabilities are given as fractions. First we investigate the impact of varying the relative probabilities of the serious and mild incidents, and the results are given in Table 1.

Table 1: Expected Survivability and $P(j)$

$a = 1.5, \pi_1 = .15, \chi_1 = .008, \pi_3 = .33, \chi_3 = .075$

cost	$P(1)=.1$	$P(1)=.5$	$P(1)=.9$
5.0	0.7712	.7340	.6964
10.00	0.8076	.7692	.7280
25.00	0.8386	.7986	.7590
50.00	0.8600	.8188	.7790
75.00	0.8756	.8372	.7994
100.00	0.8916	.8548	.8202

While the survivability increases with the cost of the defense mechanism as expected from the relation of the transition probabilities to cost, the survivability does not appear to decrease significantly with increases in the probability of occurrence of serious incidents. This is somewhat surprising, and this particular result is most likely related to the method of genera-

tion of the $p(r,s)$'s in T. This method does not vary the $p(r,s)$'s very much with j . With some other set of $\{p(r,s)\}$, we may well find greater sensitivity of survivability to the $p(j)$'s.

Table 2 shows how survivability changes with the parameter π_1 , which determines $p(1,1|m)$, the probability of remaining normal under attack given a defense mechanism m .

Table 2: Expected Survivability and π_1

$P(1)=.5, a = 1.5, \chi_1=.008, \pi_3 = .33, \chi_3 = .075$

cost	$\pi_1=.1$	$\pi_1=.15$	$\pi_1=.2$
5	0.7106	0.734	0.748
10	0.7506	0.7692	0.7786
25	0.7948	0.7986	0.8
50	0.818	0.8188	0.8188
75	0.8372	0.8372	0.8372
100	0.8548	0.8548	0.8548

Table 3 gives the survivabilities as χ_1 varies. χ_1 determines $p(1,slm)$ for $s>1$, that is, the probabilities of going to compromised states, including the nonfunctional state.

Table 3: Expected Survivability and χ_1

$P(1)=.5, a = 1.5, \pi_1 = .15, \pi_3 = .33, \chi_3 = .075$

cost	$\chi_1=.006$	$\chi_1=.008$	$\chi_1=.010$
5	0.7332	0.7342	0.734
10	0.7672	0.77	0.7692
25	0.7948	0.803	0.7986
50	0.8116	0.8288	0.8188
75	0.8244	0.8508	0.8372
100	0.8372	0.8714	0.8548

Tables 4 and 5 show the effect of varying π_3 and χ_3 respectively.

Table 4: Expected Survivability and π_3

$P(1)=.5$, $a = 1.5$, $\pi_1 = .15$, $\chi_1 = .008$, $\chi_3 = .075$

cost	$\pi_3 = .30$	$\pi_3 = .33$
5	0.7258	0.734
10	0.7588	0.7692
25	0.7926	0.7986
50	0.8118	0.8188
75	0.829	0.8372
100	0.846	0.8548

π_3 determines the levels of the transition probability $p(1,1)$ as cost changes. Thus the higher the value of π_3 , the higher the chances that the system will stay in the normal state, and thus the survivability will be higher. This is what we observe from Table 4, and we also notice that the impact is relatively greater at lower costs than at higher costs.

Table 5: Expected Survivability and χ_3

$P(1)=.5$, $a = 1.5$, $\pi_1 = .15$, $\chi_1 = .008$, $\pi_3 = .33$

cost	$\chi_3 = .080$	$\chi_3 = .075$
5	0.7288	0.734
10	0.7618	0.7692
25	0.7946	0.7986
50	0.8134	0.8188
75	0.8322	0.8372
100	0.8496	0.8548

χ_3 determines the levels of the transition probabilities $p(1,s)$ for $s > 1$, that is, the compromise probabilities. Thus a (slightly) higher value of χ_3 leads to lower values of survivability. The

relative impact is not insignificant, since the change in χ_3 is very small, and the impact is constant over the values of cost.

Another quantity of interest is the “average damage caused per unit time.” This is computed by taking the total damage (= sum of (1-surv) over all episodes) and dividing by the total time elapsed during the simulation. Thus if the rate of arrivals of incidents is increased, this simply amounts to accelerating the time scale, and expected survivability remains the same. However, the *average damage done* changes and this can be seen in Table 6.

Table 6: Average Damage and Arrival Rate a

$P(1)=.5, a = 1.5, \pi_1 = .15, \chi_1 = .008, \pi_3 = .33, \chi_3 = .075$

cost	a=1.5		a=2.0	
	EXP-SURV	AVE-DMG	EXP-SURV	AVE-DMG
5	0.734	0.417	0.734	0.556
10	0.7692	0.361	0.7692	0.482
25	0.7986	0.315	0.7986	0.42
50	0.8188	0.283	0.8188	0.378
75	0.8372	0.255	0.8372	0.339
100	0.8548	0.227	0.8548	0.302

A similar situation arises when considering the impact of a possible correlation between the arrival rates and the incident type. So far we have not assumed any such correlation, but there is evidence suggesting that it does exist. In fact, if the subsequent incident is of Type 2 (less serious) the inter-arrival time is shorter than if the incident is of Type 1 (more serious). Therefore, we included that effect in our simulation. Again, the expected survivability does not change, because the (less serious) incidents simply happen faster and the system responds in the same way. However, the *average damage done* increases, and this would have increased even more if the more serious incidents had occurred faster (instead of the way we have it). In this case, the damage done saturates with higher defense levels, and hence does not increase. The results are shown in Table 7.

Table 7: *Expected Survivability and Correlation of (a, j)*

$P(1)=.5$, $a = 1.5$, $\pi_1 = .15$, $\chi_1 = .008$, $\pi_3 = .33$, $\chi_3 = .075$

cost	a=1.5		a(j=1)=1.5; a(j=2)=2.0	
	EXP-SURV	AVE-DMG	EXP-SURV	AVE-DMG
5	0.734	0.417	0.734	0.472
10	0.7692	0.361	0.7692	0.409
25	0.7986	0.315	0.7986	0.357
50	0.8188	0.283	0.8188	0.321
75	0.8372	0.255	0.8372	0.288
100	0.8548	0.227	0.8548	0.257

The tables above show the absolute changes in expected survivability when some parameter is varied. To fully understand the impact of a parameter on expected survivability, we need to examine the relative changes. This is given by the elasticities, which give the percent changes in expected survivability when the parameters are varied by 100 percent. Thus these relative changes give a more accurate measure of the impacts. The elasticities (or η 's) are given in Table 8.

Table 8: *Relative Changes in Survivability with Respect to Parameter Values*

Parameter	$P(1)$	π_1	π_3	χ_1	χ_3
Relative change in SURV (η)	-0.062	0.125	0.603	0.201	-0.871

Table 8 confirms what we noticed before with respect to the insensitivity of survivability to $P(1)$. Survivability appears to be most sensitive to π_1 and χ_3 . That is, the level of the transition probabilities is most important, rather than how they change with m .

In the above simulations, we have assumed a Poisson process, which is a flexible model and commonly used in point processes. However, any other distribution may be used instead, and the impact of alternative distributions such as the mixed exponential may be investigated. For example, a mixture of two exponentials may be reasonable. This may arise from possible two types of attackers (amateur and experienced), each with its own rate of attacking. Then $\{f(t) = f(t; a_1, a_2, \alpha)\}$, where a_1 may be the rate for amateurs, a_2 the rate for experienced attackers, and α the proportion of amateurs to experienced attackers.

Another assumption made in this simulation is that the rate at which incidents occur is constant over time. However, there may well be a trend, and the impact of a trend in the arrival rate “a”, say $a = a_0 + a't$, also needs to be investigated.

The above results are just a small subset of all the possible analyses that can be done with this simulation model but they demonstrate the potential of this model and this approach.

Any incidents-process can be generated, and any system-response may be inserted in the model through the transition matrix T . Thus we can investigate the survivability and the damage done for any scenario for any set of defense mechanisms. Given the costs of these mechanisms, we can derive a survivability/cost function as shown below, and achieve a cost-effective level of security.

In Figure 3, we have plotted the expected survivability against cost for $P(1)=.5$ and other parameter values as in Table 1. The plot shows the relationship between cost and survivability. As cost increases, survivability increases rapidly at first, and then more slowly. Such a plot can provide a systems manager with the ability to make an informed decision about the level of defense that is most appropriate for his or her organization.

When survivability is not critical, the organization may choose a lower point on the trade-off curve, but when survivability is critical, the organization may well choose a point higher up on the curve. In the case when the “indifference curve” can be estimated, we can actually choose an optimal or “best” point on the curve. However, even if we are not aiming for optimality, we can still use the curve to find the most appropriate point in the tradeoff between cost and survivability.

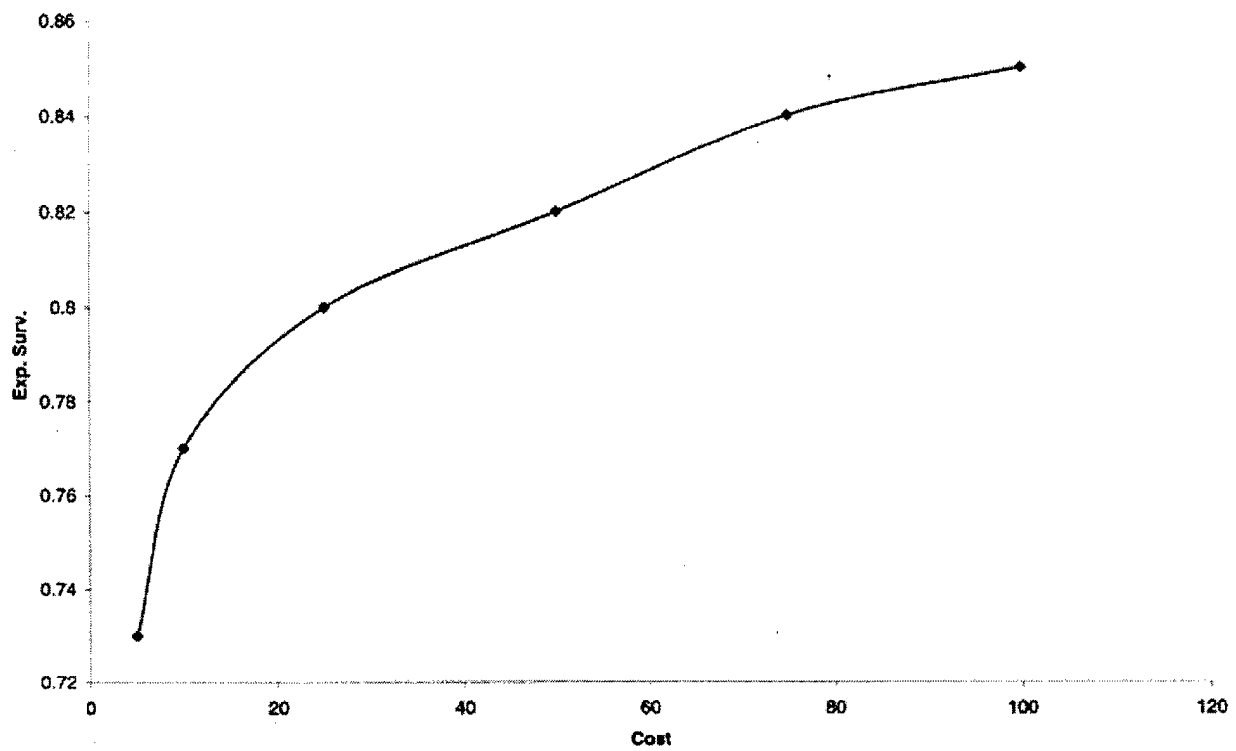


Figure 3: Expected Survivability Against Cost

5 Summary and Conclusions

5.1 Summary

In this paper we have developed a model that simulates complete episodes of attacks on network computer systems and the responses of these systems. This approach has involved developing a flexible template that can be used to analyze a variety of scenarios related to attacks on and survivability of network systems. The model encompasses several detailed aspects of the attack incidents, such as the type of attack, the number of attackers, and possible correlation between the rate of incidents and the type of incidents.

It can easily be extended to include trends or alternative distribution of inter-incident times.

The system response has been modeled probabilistically through a state transition matrix where the transition probabilities are functions of the type of incident and the defense mechanism. We have outlined a set of reasonable constraints on the probabilities and a model to generate them in the absence of data. The model reflects a relationship between the transition probabilities and the cost of the defense mechanism the system may have. The model can also be calibrated by expert opinion.

We have also proposed some survivability measures analogous to those used in telecommunications. These take into account the different functionalities and services that a system may have to perform and the different priorities they may be assigned.

We have used this model to simulate episodes with different parameter values to see the sensitivity of the results to the parameters. In each case, we have computed the expected survivability as the cost of defense varies. We have defined the concept of "average damage done per unit time," and found it useful to assess the impact on systems under various scenarios. Thus the model has been demonstrated to be capable of cost/benefit analysis that should be useful for systems managers and CIOs in managing the security of their systems. We show further that simple heuristics may be used to enhance survivability cost-effectively.

5.2 Conclusions

The main conclusions and insights may be summarized as follows:

- $P(j)$ did not have a great impact with the range of values used here. This is rather surprising, and this finding needs to be investigated with other parameter values, particularly with different values of the transition matrix.

- The occurrence rate of incidents does not have any impact of expected survivability, because it is computed *per incident*. However, the *average damage done* changes significantly. A similar result holds for a correlation between the occurrence rate and type of incident. Expected survivability does not change, but the *average damage done* changes significantly.
- Some of the parameters have significant impacts. The absolute change in the values of survivability does not appear large, but that is because the changes in the parameter values were small. The *relative* changes in survivability with respect to some of the parameter values were significant. These sensitivities are expected to be nonlinear, so many more runs need to be done when data are available. The particular results given here are simply to illustrate what can be accomplished with this model, and only relate to the hypothetical data used.
- The concept of average damage is a useful one, and should be included in future studies.
- The model developed here can help managers make better decisions to enhance the security of network systems against attack incidents. With real data, managers should be able to achieve decisions closer to optimality via the survivability/cost curves, and knowing their own tolerance for risk.

Some of the key assumptions that were made in this analysis, and which should be relaxed in the future are

- return of the system to the normal state when an incident occurs
- an arbitrary model to generate the transition probabilities
- arbitrary values for costs of defense mechanisms
- arbitrary values for survivabilities at system states
- transition probabilities constant over time, i.e., no learning

Most of these assumptions can be relaxed within the model. What are required are additional data, for example, on the transition times of systems through various states, and what those states tend to be under the different attacks.

5.3 Policy Implications

The main policy implications of this analysis are that more data should be collected so that models like these can be run with realistic parameter values. Running this model will also indicate where more information is vital and which vulnerabilities most require protection.

Our main recommendations regarding further data collection are

1. Maintain incidents-data in a standard database for easy access for analysis after sanitizing.
2. Map details of incidents into broad categories in terms of costs, impacts, and survivability.

3. For each incident, order or rank the methods of operations (MOs) by some criteria.
4. Collect data on inroads made for each incident and the "end" state of systems after an "episode."
5. Utilize more detail in notes (NT) and corrective action (CA) classifications. Trace data on perpetrators.
6. Identify attacking site and number of the attacking site.
7. Model learning on the part of attackers and on the part of victim sites.
8. Interview "attackers" anonymously or when caught.
9. Obtain more information on behavior of attackers and their motivation, to predict incidents.
10. Obtain data on various reaction and response times of victim sites.
11. Research long-term precautions that victimized sites take.
12. Obtain more information on the victim site (e.g., kind of site, defense mechanisms in place).

6 Future Work

A major need in survivability management is for more data collection (as outlined above) so that managers can better assess survivability and security and can make better decisions regarding the costs and benefits of alternative defense strategies for their systems. In particular, we need data on the response probabilities of different systems as functions of their configurations. We also need data on the intermediate states that systems may go through when attacked as well as the transition times.

With more data, further modeling and analysis could be done to explore specific systems and specific defense policies. This would have significant practical benefit to end users of information technology. Also, such a model may be embedded in a DSS that CIOs could use to make decisions.

Further simulation should be done to explore the relationship between survivability and the many parameters in the overall model, which includes the incidents process, the response matrix, and the survivability measures. Alternative models of the incidents process may also be explored. For example, there might be analogies with radiation bursts that cause damage to the atmosphere, or epidemiological models regarding the spread of disease.

Another area requiring further research is the organizations' and managers' evaluation of their own information systems. We need to know what functions and services of their systems are important to them. Thus a survey of IS managers and CIOs would be extremely useful. The survey needs to be designed very carefully to correctly elicit the tradeoffs managers would be willing to make, and how they value different aspects of their ISs. Methods like conjoint analysis may be very valuable for this. We may also need to employ decision analysis techniques to evaluate the managers' indifference curve related to cost and survivability. Integrating the survivability/cost curve with the indifference curve would lead to "optimal" or at least improved solutions. Data from the insurers of ISs may also be useful.

Finally, we need to model learning on the part of both attackers and victims. We need to know how victimized sites are upgraded with respect to defenses, and how attackers learn from their activities and outcomes of their attacks. Presumably there is a continuous cycle of increasing sophistication on both sides, and we need to incorporate this into the model.

Bibliography

- [Basawa 80] Basawa, I.V. & Prakasa Rao, B.L.S. *Statistical Inference for Stochastic Processes*. New York, NY: Academic Press, 1980.
- [Cohen 99] Cohen, F. *Simulating Cyber Attacks, Defenses, and Consequences*. Fred Cohen & Associates 1999.
- [CSI 98] Computer Security Institute. *Computer Security Issues and Trends*. 4, 1 (Winter 1998).
- [Daley 88] Daley, D.J. & Vere-Jones, D. *An Introduction to the Theory of Point Processes*. New York, NY: Springer-Verlag, 1988.
- [Ellison 97] Ellison, R.J.; Fisher, D.A.; Linger, R.C.; Lipson, H.F.; Longstaff, T. & Mead, N.R. *Survivable Network Systems: An Emerging Discipline* (CMU/SEI-97-TR-013 ADA 341963) Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 1987
Available WWW<URL <http://www.sei.cmu.edu/publications/documents/97.reports/97tr013/97tr013abstract.html>>
- [Fisher 99] Fisher, D.A. "Emergent Algorithms—A New Method for Enhancing Survivability in Unbounded Systems," *IEEE Proceedings of the Hawaii International Conference on Systems Sciences*. Wailea, HI, Jan. 5-7, 1999. New York: IEEE Computer Society Press, 1990.
- [Howard 95] Howard, J. *An Analysis of Security Incidents on the Internet (1989-1995)*. Ph.D. Dissertation, Carnegie-Mellon University, Pittsburgh, PA, 1995.
- [Howard 98] Howard, J. & Longstaff, T. *A Common Language for Computer Security Incidents*. (SAND98-8667). Livermore, CA: Sandia National Laboratories, 1998.
- [Law 82] Law, A.M. & Kelton, W.D. *Simulation Modeling and Analysis*. New York, NY: McGraw-Hill, 1982.

- [Lilien 92]** Lilien, G.L.; Kotler, P.; & Moorthy, K.S. *Marketing Models*. Englewood Cliffs, NJ: Prentice Hall, 1992.
- [Linger 98]** Linger, R.C.; Mead, N.R.; & Lipson, H.F. *Requirements Definition for Survivable Network Systems*. 1998 by IEEE. Proceedings of the International Conference on Requirements Engineering, Colorado Springs, CO: April 6-10, 1998. New York, IEEE Computer Society Press. Also published Pittsburgh, PA: Software Engineering Institute, Carnegie-Mellon University, 2000 Available WWW<URL <http://www.sei.cmu.edu/programs/nss/icre.html>>
- [Moitra 97]** Moitra, S.D.; Oki, E.; & Yamanaka, N. *Some New Survivability Measure for Network Analysis and Design*. IEICE Transactions on Communications. *E80-B*, 4, April 1997.
- [Snyder 91]** Snyder, D.L. & Miller, M.I. *Random Point Processes in Time and Space*. New York, NY: Springer-Verlag, 1991.

REPORT DOCUMENTATION PAGEForm Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (LEAVE BLANK)	2. REPORT DATE December 2000	3. REPORT TYPE AND DATES COVERED Final	
4. TITLE AND SUBTITLE A Simulation Model for Managing Survivability of Networked Information Systems		5. FUNDING NUMBERS C — F19628-95-C-0003	
6. AUTHOR(S) Soumyo D. Moitra, Suresh L. Konda			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213		8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2000-TR-021	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ ESC/XPK 5 Eglin Street Hanscom AFB, MA 01731-2116		10. SPONSORING/MONITORING AGENCY REPORT NUMBER ESC-TR-2000-020	
11. SUPPLEMENTARY NOTES			
12.A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS		12.B DISTRIBUTION CODE	
13. ABSTRACT (MAXIMUM 200 WORDS) In this paper we develop a model to evaluate the tradeoffs between the cost of defense mechanisms for networked systems and the resulting expected survivability after a network attack. The model consists of three submodels. The first submodel simulates the occurrence of attacks or incidents. The second submodel simulates the impact of an attack on the system. This depends on the type of attack and the defense mechanism installed in the system. The third submodel assesses the survivability of the system which depends on the degree of its degradation after the attack. By varying the level of defense in the simulation, we examine how this expected survivability changes with the defense level. Since costs are assumed to increase with the strength of the defense system, we can derive a cost/survivability curve that managers can use to decide on the appropriate level of security for their organizations. We have also explored the sensitivity of expected survivability to various parameters of the model, such as, the mix of attack types and the rate of occurrence of incidents.			
14. SUBJECT TERMS survivability, network systems, transition probabilities, defense mechanisms, incident types		15. NUMBER OF PAGES 44	
16. PRICE CODE			
7. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102